

RECYCLING BICG WITH AN APPLICATION TO MODEL REDUCTION*

KAPIL AHUJA[†], ERIC DE STURLER[†], SERKAN GUGERCIN[†], AND EUN R. CHANG[†]

Abstract. Science and engineering problems frequently require solving a sequence of dual linear systems. Besides having to store only few Lanczos vectors, using the BiConjugate Gradient method (BiCG) to solve dual linear systems has advantages for specific applications. For example, using BiCG to solve the dual linear systems arising in interpolatory model reduction provides a backward error formulation in the model reduction framework. Using BiCG to evaluate bilinear forms – for example, in quantum Monte Carlo (QMC) methods for electronic structure calculations – leads to a quadratic error bound. Since our focus is on sequences of dual linear systems, we introduce *recycling BiCG*, a BiCG method that recycles two Krylov subspaces from one pair of dual linear systems to the next pair. The derivation of recycling BiCG also builds the foundation for developing recycling variants of other bi-Lanczos based methods, such as CGS, BiCGSTAB, QMR, and TFQMR.

We develop an augmented bi-Lanczos algorithm and a modified two-term recurrence to include recycling in the iteration. The recycle spaces are approximate left and right invariant subspaces corresponding to the eigenvalues closest to the origin. These recycle spaces are found by solving a small generalized eigenvalue problem alongside the dual linear systems being solved in the sequence.

We test our algorithm in two application areas. First, we solve a discretized partial differential equation (PDE) of convection-diffusion type. Such a problem provides well-known test cases that are easy to test and analyze further. Second, we use recycling BiCG in the Iterative Rational Krylov Algorithm (IRKA) for interpolatory model reduction. IRKA requires solving sequences of slowly changing dual linear systems. We analyze the generated recycle spaces and show up to 70% savings in iterations. For our model reduction test problem, we show that solving the problem without recycling leads to (about) a 50% increase in runtime.

Key words. Krylov subspace recycling, deflation, bi-Lanczos method, Petrov-Galerkin formulation, BiCG, model reduction, rational Krylov, \mathcal{H}_2 approximation.

AMS subject classifications. 65F10, 65N22, 93A15, 93C05.

1. Introduction. We focus on solving the sequence of dual linear systems,

$$A^{(j)}x^{(j)} = b^{(j)}, \quad A^{(j)*}\tilde{x}^{(j)} = \tilde{b}^{(j)}, \quad (1.1)$$

where $A^{(j)} \in \mathbb{C}^{n \times n}$ and $b^{(j)}, \tilde{b}^{(j)} \in \mathbb{C}^n$ vary with j , the matrices $A^{(j)}$ are large and sparse, the solution of the dual system is relevant, and the change from a pair of systems to the next is small.

In several application areas, there are important advantages to solving dual linear systems using the BiCG algorithm [20]. BiCG has a short-term recurrence, so very few Lanczos vectors have to be stored. In addition, using BiCG to solve the dual linear systems arising in interpolatory model reduction provides a backward stable method (with respect to the interpolation conditions) for computing a reduced order model [12] (see Section 5.2). This makes BiCG attractive even for symmetric positive definite (SPD) systems. Furthermore, in several applications, such as QMC algorithms [5], we need to evaluate bilinear forms of the type $u^*A^{-1}w$, where $u, w \in \mathbb{C}^n$ and A is non-Hermitian. Solving dual linear systems for u and w to compute $u^*A^{-1}w$ provides a quadratic error bound [48].

Since BiCG is advantageous for solving dual linear systems and we need to solve a sequence of such systems, we focus on Krylov subspace recycling for BiCG. We refer

*This material is based upon work supported by the National Science Foundation under Grant No. NSF-EAR 0530643, NSF-DMS 1025327, and NSF-DMS 0645347.

[†]Department of Mathematics, Virginia Tech, Blacksburg, VA 24061.

to our recycling BiCG method as RBiCG. In addition, the BiCG algorithm forms the basis of other popular bi-Lanczos based algorithms like CGS [45], BiCGSTAB [50], QMR [24], and TFQMR [22]. Hence, the derivation of RBiCG is also useful for developing recycling variants of these algorithms [3].

The convergence of Krylov subspace methods for solving a linear system, to a great extent, depends on the spectrum of the matrix, and the deflation of eigenvalues close to the origin usually improves the convergence rate [37, 47]. If the Krylov subspace is augmented with an eigenvector, then the associated eigenvalue is effectively deflated. Likewise, for BiCG, it can be shown that if the dual Krylov subspace, $K^i(A^{(j)*}, \tilde{r}_0)$, is augmented with left eigenvectors, the corresponding right eigenvectors are removed from the primal residual (and vice versa if the primal Krylov subspace is augmented with right eigenvectors) [15]. Therefore, while solving a pair of systems, we select approximate left- and right invariant subspaces of $A^{(j)}$ (corresponding to small eigenvalues in absolute value), and use these to accelerate the solution of the next pair of systems. This process is called *Krylov subspace recycling*, and leads to faster convergence for the next pair of systems.

For solving a single linear system, ‘recycling’ has been used in the GCROT [16] and the GMRES-DR [37] algorithms. For solving a sequence of linear systems, this idea was first proposed in [38], where it is applied to the GCROT and the GCRO-DR algorithms. Recycling techniques are adapted to short term recurrences in the RMINRES [53] algorithm; see [36] for an improved version. GCROT as in [38], GCRO-DR, and RMINRES all focus on solving a sequence of single systems rather than a sequence of two dual systems, which is the focus here. For a fixed matrix with multiple right hand sides deflation-based approaches are proposed in [2, 1]. For a comprehensive discussion of recycling algorithms see [38].

In addition to testing RBiCG for IRKA [29] for interpolatory model reduction, we test RBiCG for a model convection-diffusion problem. PDEs of this type are pervasive in science and engineering, they lead to nonsymmetric matrices for which BiCG may be well-suited, and they provide well-known test cases that are easy to reproduce and to analyze further. Convection-diffusion problems arise, for example, in the Oseen problem (a fixed-point linearization of the Navier-Stokes equations), in chemically-reacting flows, heat flow in a medium with transport, and so on. Moreover, any large discretized PDE leads to a potential model reduction problem, for example, for uncertainty quantification, for optimizing an engineering process, or indirectly estimating parameters in the model using measurements. We analyze the generated recycle spaces for both test problems, and we show up to 70% reduction in the iteration count. For our model reduction test problem, using BiCG instead of RBiCG would take approximately 50% more time to generate the reduced order model. As recycling is not needed for every pair of linear systems, this means that the improvement in time for those systems where recycling is actually used is substantially larger (see section 6).

To simplify notation, we drop the superscript j in (1.1). At any particular point in the sequence of systems, we refer to $Ax = b$ as the primary system and $A^*\tilde{x} = \tilde{b}$ as the dual system. Throughout the paper, $\|\cdot\|$ refers to the two-norm, (\cdot, \cdot) refers to the standard inner product, and e_i is the i -th canonical basis vector. Unless otherwise stated, we refer to the primary system recycle space and the dual system recycle space collectively as the recycle space.

In the next section, we briefly discuss the BiCG algorithm, and in section 3 we derive the RBiCG algorithm using a previously computed recycle space. How to

compute or update such a recycle space efficiently is discussed in section 4. After explaining the basics of interpolatory model reduction, we discuss how RBiCG is applied in IRKA in section 5. We present numerical experiments and results in section 6 and conclusions in Section 7.

2. The BiCG algorithm. For the primary system, let x_0 be the initial guess with residual $r_0 = b - Ax_0$. Krylov subspace methods, in general, find approximate solutions by projection onto the Krylov subspace associated with A and r_0 [51]. The i -th solution iterate is given by

$$x_i = x_0 + \varrho_i, \quad (2.1)$$

where $\varrho_i \in K^i(A, r_0) \equiv \text{span}\{r_0, Ar_0, A^2r_0, \dots, A^{i-1}r_0\}$ is defined by some projection. The BiCG method defines this projection using the Krylov subspace associated with the dual system, leading to two bi-orthogonal bases and a pair of three-term or coupled two-term recurrences. This method is called the bi-Lanczos method [34, 20]. We initialize the Lanczos vectors as follows:

$$v_1 = r_0 / \|r_0\|, \quad \tilde{v}_1 = \tilde{r}_0 / \|\tilde{r}_0\|.$$

Defining $V_i = [v_1 \ v_2 \ \dots \ v_i]$ and $\tilde{V}_i = [\tilde{v}_1 \ \tilde{v}_2 \ \dots \ \tilde{v}_i]$, the $(i+1)$ -th Lanczos vectors are given by

$$\gamma v_{i+1} = Av_i - V_i \tau \perp \tilde{V}_i, \quad \tilde{\gamma} \tilde{v}_{i+1} = A^* \tilde{v}_i - \tilde{V}_i \tilde{\tau} \perp V_i,$$

where the scalars γ and $\tilde{\gamma}$ and the vectors τ and $\tilde{\tau}$ are to be determined. This bi-orthogonality condition leads to a pair of 3-term recurrences (see [41]), so that computation of the $(i+1)$ -th Lanczos vectors requires only the i -th and the $(i-1)$ -th Lanczos vectors. These 3-term recurrences are called the bi-Lanczos relations, and they are defined as follows:

$$\begin{aligned} AV_i &= V_{i+1} \underline{T}_i = V_i T_i + t_{i+1,i} v_{i+1} e_i^T, \\ A^* \tilde{V}_i &= \tilde{V}_{i+1} \tilde{\underline{T}}_i = \tilde{V}_i \tilde{T}_i + \tilde{t}_{i+1,i} \tilde{v}_{i+1} e_i^T, \end{aligned}$$

where T_i, \tilde{T}_i are $i \times i$ tridiagonal matrices, $t_{i+1,i}$ is the last element of the last row of $\underline{T}_i \in \mathbb{C}^{(i+1) \times i}$, and $\tilde{t}_{i+1,i}$ is the last element of the last row of $\tilde{\underline{T}}_i \in \mathbb{C}^{(i+1) \times i}$.

The next step is to find approximate solutions by projection. To exploit the efficiency of short-term recurrences in the bi-Lanczos algorithm, we use the bi-orthogonality condition to define the projection. This leads to a Petrov-Galerkin approach. Since the columns of V_i form a basis for $K^i(A, r_0)$, we can define ϱ_i in (2.1) as $\varrho_i = V_i y_i$, and the bi-orthogonality (or Petrov-Galerkin) condition then implies

$$r_i = b - A(x_0 + \varrho_i) = r_0 - AV_i y_i \perp \tilde{V}_i.$$

The vector y_i is defined by this orthogonality condition. The solution iterate for the dual system, \tilde{x}_i , is similarly defined by $\tilde{x}_i = \tilde{x}_0 + \tilde{V}_i \tilde{y}_i$ and $\tilde{r}_i \perp V_i$. Further simplifications lead to the standard BiCG algorithm (Algorithm 1) [20, 51].

Next, we briefly discuss the breakdown conditions in BiCG and their remedies [26, 51]. The first breakdown happens when, at any step i , $\tilde{r}_i^* r_i = 0$. This is a breakdown in the underlying bi-Lanczos algorithm and is referred to as a *serious breakdown*. There exist so-called look-ahead strategies [23, 30] to avoid this breakdown. In addition, the two-term recurrence for the solution update requires a pivotless

Algorithm 1. *BiCG (adapted from [51])*

1. Choose initial guesses x_0 and \tilde{x}_0 . Compute $r_0 = b - Ax_0$ and $\tilde{r}_0 = \tilde{b} - A^*\tilde{x}_0$.
2. **if** $(r_0, \tilde{r}_0) = 0$ **then** initialize \tilde{x}_0 to a random vector.
3. Set $p_0 = 0$, $\tilde{p}_0 = 0$, and $\beta_0 = 0$. Choose **tol** and **max_itn**.
4. **for** $i = 1 \dots \text{max_itn}$ **do**
 - ◇ $p_i = r_{i-1} + \beta_{i-1}p_{i-1}$.
 - ◇ $\tilde{p}_i = \tilde{r}_{i-1} + \beta_{i-1}\tilde{p}_{i-1}$.
 - ◇ $q_i = Ap_i$.
 - ◇ $\tilde{q}_i = A^*\tilde{p}_i$.
 - ◇ $\alpha_i = (\tilde{r}_{i-1}, r_{i-1})/(\tilde{p}_i, q_i)$.
 - ◇ $x_i = x_{i-1} + \alpha_i p_i$.
 - ◇ $\tilde{x}_i = \tilde{x}_{i-1} + \tilde{\alpha}_i \tilde{p}_i$.
 - ◇ $r_i = r_{i-1} - \alpha_i q_i$.
 - ◇ $\tilde{r}_i = \tilde{r}_{i-1} - \tilde{\alpha}_i \tilde{q}_i$.
 - ◇ **if** $\|r_i\| \leq \text{tol}$ and $\|\tilde{r}_i\| \leq \text{tol}$ **then break**.
 - ◇ $\beta_i = (\tilde{r}_i, r_i)/(\tilde{r}_{i-1}, r_{i-1})$.
5. **end for**.

LDU decomposition of the tridiagonal matrix T_i , which may not always exist. This breakdown is referred to as a breakdown of the *second kind*, and it can be avoided by performing the LDU decomposition with 2×2 block diagonal elements [8]. The breakdown conditions in RBiCG are the same, and similar solutions can be applied. Therefore, and for the sake of brevity, we do not discuss breakdowns for RBiCG separately, and we'll assume henceforth in our derivations that breakdowns do not occur. Note that extensive experiments show that BiCG works well, and that breakdowns rarely happen in practice [41, 30].

3. Recycling BiCG: Using a Recycle Space. In this section, we modify the BiCG algorithm to use a given recycle space. First, we briefly describe the recycling idea used in the GCRO-DR algorithm. After solving the j -th primary system in (1.1), GCRO-DR computes the matrices $U, C \in \mathbb{C}^{n \times k}$, such that $\text{range}(U)$ is an approximate invariant subspace of $A^{(j)}$, $A^{(j+1)}U = C$ and $C^*C = I$. It then computes an orthogonal basis for the Krylov subspace $K^i((I - CC^*)A, (I - CC^*)r_0)$. This produces the Arnoldi relation

$$AV_i = CC^*AV_i + V_{i+1}\underline{H}_i \iff (I - CC^*)AV_i = V_{i+1}\underline{H}_i,$$

where \underline{H}_i is an $(i+1) \times i$ upper Hessenberg matrix. GCRO-DR finds the residual-minimizing solution over the (direct) sum of the recycle space, $\text{range}(U)$, and the new search space generated, $\text{range}(V_i)$.

In RBiCG, we use the matrix U , derived from an approximate right invariant subspace of $A^{(j)}$, to define the primary system recycle space, and compute $C = A^{(j+1)}U$. Similarly, we use the matrix \tilde{U} , derived from an approximate left invariant subspace of $A^{(j)}$, to define the dual system recycle space, and compute $\tilde{C} = A^{(j+1)*}\tilde{U}$. Instead of C being an orthogonal matrix, U and \tilde{U} are computed such that C and \tilde{C} are bi-orthogonal; see Section 4.3. The number of vectors selected for recycling is denoted by k , and hence, U, \tilde{U}, C , and $\tilde{C} \in \mathbb{C}^{n \times k}$. Next, we derive an augmented bi-Lanczos algorithm that computes bi-orthogonal bases for the primal and dual Krylov subspaces. The two-term recurrence for the solution update in RBiCG is derived in Section 3.2.

3.1. The Augmented Bi-Lanczos Algorithm. The standard bi-Lanczos algorithm computes columns of V_i and \tilde{V}_i such that, in exact arithmetic, $V_i \perp_b \tilde{V}_i$, where \perp_b denotes bi-orthogonality; this implies that $\tilde{V}_i^* V_i$ is a diagonal matrix. Since we recycle spaces U and \tilde{U} , the bi-Lanczos algorithm must be modified to compute the columns of V_i and \tilde{V}_i such that either

$$[U \ V_i] \perp_b [\tilde{U} \ \tilde{V}_i] \quad (3.1)$$

or

$$[C \ V_i] \perp_b [\tilde{C} \ \tilde{V}_i]. \quad (3.2)$$

We choose to implement (3.2), because it leads to simpler algebra and hence a more efficient algorithm. It also has the advantage that the RBiCG algorithm has a form similar to the standard BiCG algorithm. Next, we derive the recurrences that implement (3.2), where $C \perp_b \tilde{C}$ has already been satisfied. The latter relation is easy to implement when computing the recycle space. Indeed, we can compute C and \tilde{C} such that $\tilde{C}^* C$ is a real, positive, diagonal matrix; see Section 4.3. As in the BiCG algorithm, we assume v_1 and \tilde{v}_1 are available from the initial residuals r_0 and \tilde{r}_0 . We make this statement more precise below. The $(i+1)$ -th Lanczos vector for the primary system is computed by

$$\gamma v_{i+1} = Av_i - V_i \tau - C \rho \perp [\tilde{C} \ \tilde{V}_i], \quad (3.3)$$

where γ , τ , and ρ are to be determined. Combining (3.2) and (3.3), we get the following equations,

$$\begin{aligned} \mathcal{D}_c \rho &= \tilde{C}^* A v_i, \\ \mathcal{D}_i \tau &= \tilde{V}_i^* A v_i, \end{aligned} \quad (3.4)$$

where $\mathcal{D}_i = \tilde{V}_i^* V_i$ and $\mathcal{D}_c = \tilde{C}^* C$ are both diagonal matrices and \mathcal{D}_c has real, positive coefficients (see Section 4.3). As discussed before, a breakdown in the standard BiCG algorithm because of singular \mathcal{D}_i can be fixed with look-ahead strategies. Assuming breakdowns do not occur, we can solve for τ and ρ in (3.4) and choose a normalization γ ; substituting these into (3.3) gives the $(i+1)$ -th Lanczos vector. Because of the bi-orthogonality condition (3.2), the full recurrence for v_{i+1} reduces to a $(3+k)$ -term recurrence, where k is the number of columns of C . This implies that the computation of the $(i+1)$ -th Lanczos vector requires the i -th and $(i-1)$ -th Lanczos vectors and C . Similarly, we get a $(3+k)$ -term recurrence for computing the Lanczos vectors for the dual system. We refer to this pair of $(3+k)$ -term recurrences as the augmented bi-Lanczos relations; they are given by

$$(I - C\tilde{C}^*)AV_i = V_{i+1}\underline{T}_i, \quad (I - \tilde{C}\tilde{C}^*)A^*\tilde{V}_i = \tilde{V}_{i+1}\underline{\tilde{T}}_i, \quad (3.5)$$

where

$$\begin{aligned} \hat{C} &= \begin{bmatrix} \frac{\tilde{c}_1}{c_1^* \tilde{c}_1} & \frac{\tilde{c}_2}{c_2^* \tilde{c}_2} & \cdots & \frac{\tilde{c}_k}{c_k^* \tilde{c}_k} \end{bmatrix} = \tilde{C} \mathcal{D}_c^{-*} = \tilde{C} \mathcal{D}_c^{-1}, \\ \check{C} &= \begin{bmatrix} \frac{c_1}{\tilde{c}_1^* c_1} & \frac{c_2}{\tilde{c}_2^* c_2} & \cdots & \frac{c_k}{\tilde{c}_k^* c_k} \end{bmatrix} = C \mathcal{D}_c^{-1}. \end{aligned} \quad (3.6)$$

Using (3.2), we can rewrite (3.5) as

$$\begin{aligned} A_1 V_i &= V_{i+1} \underline{T}_i, & \text{where} & \quad A_1 = (I - C \mathcal{D}_c^{-1} \tilde{C}^*) A (I - C \mathcal{D}_c^{-1} \tilde{C}^*), \\ A_1^* \tilde{V}_i &= \tilde{V}_{i+1} \underline{\tilde{T}}_i, & \text{where} & \quad A_1^* = (I - \tilde{C} \mathcal{D}_c^{-*} C^*) A^* (I - \tilde{C} \mathcal{D}_c^{-*} C^*), \end{aligned} \quad (3.7)$$

since $\tilde{C}^* V_i = 0$ and $C^* \tilde{V}_i = 0$. This new form of the augmented bi-Lanczos relations simplifies the derivation of the recurrence for the RBiCG solution update, because the operators (3.7) are each other's conjugate transpose. Note that the additional orthogonalizations in (3.7) need not be carried out in an actual algorithm (see Algorithm 3.2).

3.2. The Solution Update for the Augmented Bi-Lanczos Recurrence.

The i -th solution update in the RBiCG algorithm becomes

$$x_i = x_0 + U z_i + V_i y_i, \quad \tilde{x}_i = \tilde{x}_0 + \tilde{U} \tilde{z}_i + \tilde{V}_i \tilde{y}_i, \quad (3.8)$$

With recycling, the bi-orthogonality condition (3.2) defines the Petrov-Galerkin condition,

$$r_i = r_0 - AU z_i - AV_i y_i \perp [\tilde{C} \tilde{V}_i], \quad \tilde{r}_i = \tilde{r}_0 - A^* \tilde{U} \tilde{z}_i - A^* \tilde{V}_i \tilde{y}_i \perp [C V_i]. \quad (3.9)$$

For the remainder of this section, we focus on the primary system. The derivations for the dual system are analogous. The computation of z_i and y_i can be implemented more efficiently than (3.9) suggests. Defining $\zeta = \|(I - C\hat{C}^*)r_0\|$ and $v_1 = \zeta^{-1}(I - C\hat{C}^*)r_0$, we get

$$r_0 = C\hat{C}^* r_0 + (I - C\hat{C}^*) r_0 = [C \ V_{i+1}] \begin{bmatrix} \hat{C}^* r_0 \\ \zeta e_1 \end{bmatrix}. \quad (3.10)$$

Using the augmented bi-Lanczos relation (3.5) we get

$$A [U \ V_i] \begin{bmatrix} z_i \\ y_i \end{bmatrix} = [C \ V_{i+1}] \begin{bmatrix} I & \hat{C}^* A V_i \\ 0 & \underline{T}_i \end{bmatrix} \begin{bmatrix} z_i \\ y_i \end{bmatrix}. \quad (3.11)$$

Substituting (3.10) and (3.11) in (3.9) gives

$$\begin{bmatrix} \tilde{C}^* \\ \tilde{V}_i^* \end{bmatrix} [C \ V_{i+1}] \left(\begin{bmatrix} \hat{C}^* r_0 \\ \zeta e_1 \end{bmatrix} - \begin{bmatrix} I & \hat{C}^* A V_i \\ 0 & \underline{T}_i \end{bmatrix} \begin{bmatrix} z_i \\ y_i \end{bmatrix} \right) = 0. \quad (3.12)$$

Using the bi-orthogonality condition (3.2) in the above equation we get¹

$$\begin{bmatrix} \hat{C}^* r_0 \\ \zeta e_1 \end{bmatrix} - \begin{bmatrix} I & \hat{C}^* A V_i \\ 0 & T_i \end{bmatrix} \begin{bmatrix} z_i \\ y_i \end{bmatrix} = 0. \quad (3.13)$$

Therefore, y_i and z_i are given by

$$\begin{aligned} T_i y_i &= \zeta e_1, \\ z_i &= \hat{C}^* r_0 - \hat{C}^* A V_i y_i. \end{aligned} \quad (3.14)$$

Substituting (3.14) in (3.8) leads to the following solution update:

$$x_i = x_0 + U \hat{C}^* r_0 + (I - U \hat{C}^* A) V_i y_i,$$

where y_i is obtained from solving $T_i y_i = \zeta e_1$. All computations here are done with matrix-vector products and $U \hat{C}^* A$ is not computed explicitly.

¹Note that the length of the vector e_1 in (3.13) is one less than that of e_1 in (3.12), although both denote the first canonical basis vector. Also, T_i in (3.13) is \underline{T}_i without the last row, and hence is an $i \times i$ tridiagonal matrix.

Algorithm 2. *RBICG*

1. Given U and \tilde{U} compute \hat{C} and \hat{C} using (3.6). If U and \tilde{U} are not available, then initialize U , \tilde{U} , \hat{C} , and \hat{C} to empty matrices.
2. Choose x_{-1} , \tilde{x}_{-1} and compute x_0 , \tilde{x}_0 , r_0 , and \tilde{r}_0 using (3.15).
3. **if** $(r_0, \tilde{r}_0) = 0$ **then** initialize \tilde{x}_{-1} to a random vector.
4. Set $p_0 = 0$, $\tilde{p}_0 = 0$, and $\beta_0 = 0$. Choose **tol** and **max_itn**.
5. **for** $i = 1 \dots \text{max_itn}$ **do**
 - ◇ $p_i = r_{i-1} + \beta_{i-1}p_{i-1}$; $\tilde{p}_i = \tilde{r}_{i-1} + \tilde{\beta}_{i-1}\tilde{p}_{i-1}$
 - ◇ $z_i = Ap_i$; $\tilde{z}_i = A^*\tilde{p}_i$
 - ◇ $\zeta_i = \hat{C}^*z_i$; $\tilde{\zeta}_i = \tilde{C}^*\tilde{z}_i$
 - ◇ $q_i = z_i - C\zeta_i$; $\tilde{q}_i = \tilde{z}_i - \tilde{C}\tilde{\zeta}_i$
 - ◇ $\alpha_i = (\tilde{r}_{i-1}, r_{i-1})/(\tilde{p}_i, q_i)$; $\tilde{\alpha}_i = \tilde{\alpha}_i$
 - ◇ $\zeta_c = \zeta_c + \alpha_i\zeta_i$; $\tilde{\zeta}_c = \tilde{\zeta}_c + \tilde{\alpha}_i\tilde{\zeta}_i$
 - ◇ $x_i = x_{i-1} + \alpha_ip_i$ $\tilde{x}_i = \tilde{x}_{i-1} + \tilde{\alpha}_i\tilde{p}_i$
 - ◇ $r_i = r_{i-1} - \alpha_iq_i$ $\tilde{r}_i = \tilde{r}_{i-1} - \tilde{\alpha}_i\tilde{q}_i$
 - ◇ **if** $\|r_i\| \leq \text{tol}$ and $\|\tilde{r}_i\| \leq \text{tol}$ **then break**
 - ◇ $\beta_i = (\tilde{r}_i, r_i)/(\tilde{r}_{i-1}, r_{i-1})$
6. **end for**
7. $x_i = x_i - U\zeta_c$; $\tilde{x}_i = \tilde{x}_i - \tilde{U}\tilde{\zeta}_c$

We introduce a slight change of notation to make future derivations simpler. Let x_{-1} and \tilde{x}_{-1} be the initial guesses and $r_{-1} = b - Ax_{-1}$ and $\tilde{r}_{-1} = \tilde{b} - A^*\tilde{x}_{-1}$ the corresponding initial residuals. We define

$$\begin{aligned} x_0 &= x_{-1} + U\hat{C}^*r_{-1}, & r_0 &= (I - C\hat{C}^*)r_{-1}, \\ \tilde{x}_0 &= \tilde{x}_{-1} + \tilde{U}\tilde{C}^*\tilde{r}_{-1}, & \tilde{r}_0 &= (I - \tilde{C}\tilde{C}^*)\tilde{r}_{-1}, \end{aligned} \quad (3.15)$$

and follow this convention for x_0 , \tilde{x}_0 , r_0 , and \tilde{r}_0 for the rest of the paper. Let

$$\begin{aligned} T_i &= L_i D_i R_i, \\ G_i &= (I - U\hat{C}^*A)V_i R_i^{-1}, \\ \varphi_i &= \zeta D_i^{-1} L_i^{-1} e_1. \end{aligned}$$

As in the standard BiCG algorithm, an LDU decomposition (without pivoting) of T_i might not always exist. We can avoid this breakdown in the same way as done for BiCG (see Section 2). The two-term recurrence for the solution update of the primary system is now given by

$$x_i = x_{i-1} + \varphi_{i,i} G_i e_i \quad \text{for } i \geq 1,$$

where $\varphi_{i,i}$ is the last entry of the vector φ_i , and x_0 is given by (3.15). An analogous update can be derived for the dual system. Note that we never compute any explicit matrix inverse. The matrices under consideration, D_i , L_i , and R_i , are diagonal, lower triangular, and upper triangular respectively.

This two-term recurrence can be simplified such that T_i is not needed explicitly. To derive further simplifications, we use the operator A_1 (instead of A) and follow steps similar to the ones used in the derivation of BiCG [30]. Algorithm 2 provides an outline of RBICG. Some algorithmic improvements to make the code faster are not given here; see [3] for further details.

4. Recycling BiCG: Computing a Recycle Space. We use the matrices U and \tilde{U} to define the primary and dual system recycle spaces. The recycle space used in solving a linear system is fixed throughout the RBiCG iteration; however, the basis of the recycle space for the next pair of linear systems is updated periodically using the bi-Lanczos vectors. We use harmonic Ritz vectors, with respect to the current Krylov subspace, to approximate left- and right invariant subspaces cheaply.

We use the following definition [44]. Let S be a subspace of \mathbb{C}^n . Then $\lambda \in \mathbb{C}$ is a harmonic Ritz value of A and $u \in S \neq 0$ its corresponding harmonic Ritz vector with respect to the subspace $\mathcal{W} = AS$ if

$$(Au - \lambda u) \perp AS. \quad (4.1)$$

In Section 4.1, we derive a small generalized eigenvalue problem whose solution gives the desired approximate invariant subspace. The first pair of systems in our sequence of dual linear systems requires special attention, since there is no recycle space available at the start. We discuss this case in Section 4.2. In Section 4.3, we describe the construction of the bi-orthogonal C and \tilde{C} in (3.2) such that $\mathcal{D}_c = \tilde{C}^* C$ has positive real coefficients. Although, the generalized eigenvalue problem derived in Section 4.1 is of a small dimension, it would be expensive to set up in a straightforward manner. We show how to set up the problem efficiently using recurrences in Section 4.4.

4.1. Computing an Approximate Invariant Subspace. We need a sequence of consecutive Lanczos vectors v_i and \tilde{v}_i and tridiagonal matrices T_i and \tilde{T}_i to build the recycle space. There is a degree of freedom in choosing the scaling of the Lanczos vectors [26, 30, 41]. The following scaling yields $\tilde{T}_i = T_i^*$ (using (3.7) and (4.2)):

$$\|v_i\| = 1, \quad (v_i, \tilde{v}_i) = 1. \quad (4.2)$$

Hence, the Lanczos vectors are computed as follows:

$$v_i = \frac{r_{i-1}}{\|r_{i-1}\|}, \quad \tilde{v}_i = \frac{\tilde{r}_{i-1}}{(v_i, \tilde{r}_{i-1})}.$$

T_i can be computed using the residuals and iteration scalars of the RBiCG iteration as follows [2, 3]:

$$T_i = \begin{pmatrix} \frac{1}{\alpha_1} & -\frac{\|r_0\|}{\|r_1\|} \cdot \frac{\beta_1}{\alpha_1} & & & \\ -\frac{\|r_1\|}{\|r_0\|} \cdot \frac{1}{\alpha_1} & \frac{1}{\alpha_2} + \frac{\beta_1}{\alpha_1} & -\frac{\|r_1\|}{\|r_2\|} \cdot \frac{\beta_2}{\alpha_2} & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & -\frac{\|r_{i-1}\|}{\|r_{i-2}\|} \cdot \frac{1}{\alpha_{i-1}} & -\frac{\|r_{i-2}\|}{\|r_{i-1}\|} \cdot \frac{\beta_{i-1}}{\alpha_{i-1}} \end{pmatrix}.$$

Instead of using all the Lanczos vectors to update the recycle space, we update the recycle space periodically. This strategy keeps the memory requirements modest [53], as it allows us to discard Lanczos vectors periodically. The iteration process between two updates of the recycle space is referred to as a “cycle”. The length of the cycle, s , refers to the number of iterations between updates. Let V_j and \tilde{V}_j contain the Lanczos vectors generated during the j^{th} cycle,

$$V_j = [v_{(j-1)s+1} \quad \dots \quad v_{js}], \quad \tilde{V}_j = [\tilde{v}_{(j-1)s+1} \quad \dots \quad \tilde{v}_{js}].$$

Also, let

$$\Upsilon_j = \begin{bmatrix} v_{(j-1)s} & V_j & v_{js+1} \end{bmatrix}, \quad \tilde{\Upsilon}_j = \begin{bmatrix} \tilde{v}_{(j-1)s} & \tilde{V}_j & \tilde{v}_{js+1} \end{bmatrix},$$

where $v_{(j-1)s}$ and $\tilde{v}_{(j-1)s}$ are the last Lanczos vectors from the previous cycle, and v_{js+1} and \tilde{v}_{js+1} are the first Lanczos vectors from the next cycle. The augmented bi-Lanczos relations for the j^{th} cycle are now given by

$$(I - C\hat{C}^*)AV_j = \Upsilon_j\Gamma_j, \quad (I - \tilde{C}\tilde{C}^*)A^*\tilde{V}_j = \tilde{\Upsilon}_j\tilde{\Gamma}_j, \quad (4.3)$$

where $\Gamma_j, \tilde{\Gamma}_j \in \mathbb{C}^{(s+2) \times s}$ are T_j, \tilde{T}_j , respectively, with an extra row at the top (corresponding to $v_{(j-1)s}$ and $\tilde{v}_{(j-1)s}$) and at the bottom (corresponding to v_{js+1} and \tilde{v}_{js+1}).

The discussion in this paragraph concerns only the primary system. However, an analogous discussion applies to the dual system. Let U define the recycle space available from the previous linear system and U_{j-1} the recycle space generated at the end of cycle $(j-1)$ for the current linear system. We want to obtain an improved U_j from V_j, U_{j-1} , and U . It is important to note that U_j is not used for solving the current linear system. At the end of solving the current linear system, the final U_j will be U for the next linear system. There are several choices for selecting U_j [53]. For simplicity, we build U_j from $\text{range}([U_{j-1} \ V_j])$.

Based on the choices discussed in the previous two paragraphs, we first define certain matrices, and then we derive the generalized eigenvalue problem whose solution gives the approximate invariant subspace. Let

$$\begin{aligned} \Phi_j &= [U_{j-1} \ V_j], & \Psi_j &= [C \ C_{j-1} \ \Upsilon_j], & H_j &= \begin{bmatrix} 0 & B_j \\ I & 0 \\ 0 & \Gamma_j \end{bmatrix}, \\ \tilde{\Phi}_j &= [\tilde{U}_{j-1} \ \tilde{V}_j], & \tilde{\Psi}_j &= [\tilde{C} \ \tilde{C}_{j-1} \ \tilde{\Upsilon}_j], & \tilde{H}_j &= \begin{bmatrix} 0 & \tilde{B}_j \\ I & 0 \\ 0 & \tilde{\Gamma}_j \end{bmatrix}, \end{aligned}$$

where $C_{j-1} = AU_{j-1}$, $B_j = \hat{C}^*AV_j$, $\tilde{C}_{j-1} = A^*\tilde{U}_{j-1}$, and $\tilde{B}_j = \tilde{C}^*A^*\tilde{V}_j$. Then, the augmented bi-Lanczos relations (4.3) lead to

$$A\Phi_j = \Psi_j H_j, \quad A^*\tilde{\Phi}_j = \tilde{\Psi}_j \tilde{H}_j. \quad (4.4)$$

In RMINRES [53], harmonic Ritz pairs of A with respect to the subspace $\text{range}(A\Phi_j)$ have been successfully used to build the recycle space. Since we work in a Petrov-Galerkin framework, it is more intuitive to use harmonic Ritz pairs with respect to the subspace $\text{range}(A^*\tilde{\Phi}_j)$, following [9]. This leads to simpler algebra and cheaper computations. Let (λ, u) denote an harmonic Ritz pair of A . Then, we derive λ and $u \in \text{range}(\Phi_j)$ from the condition

$$(Au - \lambda u) \perp \text{range}(A^*\tilde{\Phi}_j). \quad (4.5)$$

Taking $u = \Phi_j w$ and substituting (4.4) in (4.5) gives

$$(A^*\tilde{\Phi}_j)^* A\Phi_j w = \lambda (A^*\tilde{\Phi}_j)^* \Phi_j w \Leftrightarrow (\tilde{\Psi}_j \tilde{H}_j)^* \Psi_j H_j w = \lambda (\tilde{\Psi}_j \tilde{H}_j)^* \Phi_j w.$$

Thus, condition (4.5) leads to the generalized eigenvalue problem,

$$\tilde{H}_j^* \tilde{\Psi}_j^* \Psi_j H_j w = \lambda \tilde{H}_j^* \tilde{\Psi}_j^* \Phi_j w. \quad (4.6)$$

Let the columns of W_j be the k right eigenvectors corresponding to the eigenvalues closest to the origin. Then, we take $U_j = \Phi_j W_j$. See [3] for an analogous derivation of the dual system recycle space.

4.2. The First Linear System and the First Cycle. For the first cycle of the first system, the matrices U , \tilde{U} , U_{j-1} , and \tilde{U}_{j-1} are not available. Letting T_1 and \tilde{T}_1 denote the tridiagonal matrices for the first cycle, we consider the following eigenvalue problems:

$$T_1 w = \lambda w, \quad \tilde{T}_1 \tilde{w} = \mu \tilde{w}.$$

Since $\tilde{T}_1 = T_1^*$, we solve for the left and the right eigenvectors of T_1 , \tilde{W}_1 and W_1 respectively. Hence, we take

$$U_1 = V_1 W_1, \quad \tilde{U}_1 = \tilde{V}_1 \tilde{W}_1.$$

During the second and subsequent cycles of the first linear system, U_{j-1} and \tilde{U}_{j-1} are available, but C and \tilde{C} are not. Redefining $\Psi_j, \tilde{\Psi}_j, H_j$, and \tilde{H}_j , we get the generalized eigenvalue problem (4.6) with

$$\begin{aligned} \Phi_j &= [U_{j-1} \quad V_j], \quad \Psi_j = [C_{j-1} \quad \Upsilon_j], \quad H_j = \begin{bmatrix} I & 0 \\ 0 & \Gamma_j \end{bmatrix}, \\ \tilde{\Phi}_j &= [\tilde{U}_{j-1} \quad \tilde{V}_j], \quad \tilde{\Psi}_j = [\tilde{C}_{j-1} \quad \tilde{\Upsilon}_j], \quad \tilde{H}_j = \begin{bmatrix} I & 0 \\ 0 & \tilde{\Gamma}_j \end{bmatrix}. \end{aligned}$$

For the first cycle of each of the subsequent linear systems (i.e. $j = 1$), C and \tilde{C} are available, while U_{j-1} and \tilde{U}_{j-1} are not. Redefining $\Phi_1, \tilde{\Phi}_1, \Psi_1, \tilde{\Psi}_1, H_1$, and \tilde{H}_1 , we get the generalized eigenvalue problem (4.6) with

$$\begin{aligned} \Phi_1 &= [U \quad V_1], \quad \Psi_1 = [C \quad \underline{V}_1], \quad H_1 = \begin{bmatrix} I & B_1 \\ 0 & \underline{T}_1 \end{bmatrix}, \\ \tilde{\Phi}_1 &= [\tilde{U} \quad \tilde{V}_1], \quad \tilde{\Psi}_1 = [\tilde{C} \quad \underline{\tilde{V}}_1], \quad \tilde{H}_1 = \begin{bmatrix} I & \tilde{B}_1 \\ 0 & \underline{\tilde{T}}_1 \end{bmatrix}, \end{aligned}$$

where \underline{V}_1 and $\underline{\tilde{V}}_1$ denote $[V_1 \ v_{s+1}]$ and $[\tilde{V}_1 \ \tilde{v}_{s+1}]$ respectively.

4.3. Constructing Bi-orthogonal C_j , \tilde{C}_j and C , \tilde{C} . We need to compute the matrices C_j and \tilde{C}_j such that $C_j \perp_b \tilde{C}_j$ at the end of each cycle. After solving the generalized eigenvalue problem (4.6), we set (as initial choice) $U_j = \Phi_j W_j$, $\tilde{U}_j = \tilde{\Phi}_j \tilde{W}_j$, $C_j = AU_j$, and $\tilde{C}_j = A^* \tilde{U}_j$, and we compute the SVD

$$\tilde{C}_j^* C_j = M_j \Sigma_j N_j^*, \quad (4.7)$$

such that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_k \geq 0$. Given some tolerance $\text{tol} > 0$, we pick p such that $\sigma_p \geq \text{tol} > \sigma_{p+1}$ (with both $p = k$ and $p = 0$ possible), and redefine $M_j = [m_1, \dots, m_p]$ and $N_j = [n_1, \dots, n_p]$, where m_i and n_i are the left and right singular vectors corresponding to σ_i . Next, we redefine

$$\begin{aligned} U_j &= \Phi_j W_j N_j = [U_{j-1} \quad V_j] W_j N_j, \quad \tilde{U}_j = \tilde{\Phi}_j \tilde{W}_j M_j = [\tilde{U}_{j-1} \quad \tilde{V}_j] \tilde{W}_j M_j, \\ C_j &= AU_j = A[U_{j-1} \quad V_j] W_j N_j, \quad \tilde{C}_j = A^* \tilde{U}_j = A^* [\tilde{U}_{j-1} \quad \tilde{V}_j] \tilde{W}_j M_j. \end{aligned} \quad (4.8)$$

By construction $\tilde{C}_j^* C_j$ is diagonal with real, positive coefficients.²

Analogous to the above, at the start of each linear system (after the first), we need to compute C and \tilde{C} such that $C \perp_b \tilde{C}$, cf. (3.2), and the diagonal matrix $\mathcal{D}_c = \tilde{C}^* C$ has real, positive coefficients. Taking initially for U and \tilde{U} the final matrices U_j and \tilde{U}_j from the previous pair of linear systems, we compute $C = AU$, $\tilde{C} = A^* \tilde{U}$, and compute the SVD $\tilde{C}^* C = M \Sigma N^*$. After this we proceed as for the computation of C_j and \tilde{C}_j .

4.4. Efficiently Setting up the Generalized Eigenvalue Problem. The main cost of setting up the generalized eigenvalue problem (4.6) is in computing the matrices

$$\begin{aligned} \tilde{\Psi}_j^* \Psi_j &= \begin{bmatrix} \tilde{C}^* \\ \tilde{C}_{j-1}^* \\ \Upsilon_j^* \end{bmatrix} \begin{bmatrix} C & C_{j-1} & \Upsilon_j \end{bmatrix} = \begin{bmatrix} \mathcal{D}_c & \tilde{C}^* C_{j-1} & 0 \\ \tilde{C}_{j-1}^* C & \Sigma_{j-1} & \tilde{C}_{j-1}^* \Upsilon_j \\ 0 & \tilde{\Upsilon}_j^* C_{j-1} & \underline{I} \end{bmatrix}, \\ \tilde{\Psi}_j^* \Phi_j &= \begin{bmatrix} \tilde{C}^* \\ \tilde{C}_{j-1}^* \\ \Upsilon_j^* \end{bmatrix} \begin{bmatrix} U_{j-1} & V_j \end{bmatrix} = \begin{bmatrix} \tilde{C}^* U_{j-1} & 0 \\ \tilde{C}_{j-1}^* U_{j-1} & \tilde{C}_{j-1}^* V_j \\ \tilde{\Upsilon}_j^* U_{j-1} & \underline{\tilde{I}} \end{bmatrix}, \end{aligned}$$

where $\underline{\tilde{I}}$ is the $s \times s$ identity matrix with an extra row of zeros at the top and at the bottom. Most of the blocks in these matrices can be constructed efficiently by exploiting recurrences and various algebraic relations.

The bi-orthogonality condition (3.2) and the construction of C_j and \tilde{C}_j (4.7)-(4.8) give the following orthogonality relations.

$$\tilde{C}_{j-2} \perp \Upsilon_j, \quad C_{j-2} \perp \tilde{\Upsilon}_j. \quad (4.9)$$

Next, going from top-to-bottom and left-to-right, we analyze each block of $\tilde{\Psi}_j^* \Psi_j$ and $\tilde{\Psi}_j^* \Phi_j$ in terms of its defining recurrences and simplify it using (3.2), (4.3), (4.7), (4.8), and (4.9). Blocks whose efficient computation is obvious or has already been detailed are skipped, and we focus on computations that are at least $O(n)$.

- $\tilde{C}^* C_{j-1} = \tilde{C}^* A U_{j-1} = \begin{bmatrix} \tilde{C}^* A U_{j-2} & \tilde{C}^* A V_{j-1} \end{bmatrix} W_{j-1} N_{j-1}$
 $= \begin{bmatrix} \tilde{C}^* C_{j-2} & \tilde{C}^* (C \hat{C}^* A V_{j-1} + \Upsilon_{j-1} \Gamma_{j-1}) \end{bmatrix} W_{j-1} N_{j-1}$
 $= \begin{bmatrix} \tilde{C}^* C_{j-2} & \mathcal{D}_c B_{j-1} \end{bmatrix} W_{j-1} N_{j-1},$

For this first block, we describe its efficient computation in some detail. The cost of computing $\tilde{C}^* C_{j-1}$ by direct multiplication is $O(k^2 n)$; so, it would be expensive. However, the submatrix $\tilde{C}^* C_{j-2}$ is available from the previous cycle, and \mathcal{D}_c is a diagonal matrix independent of the cycle (so both must be computed at most once per linear system). Furthermore, B_{j-1} has been computed during the (augmented) bi-Lanczos iteration. Finally, the matrix-matrix product $[\tilde{C}^* C_{j-2} \quad \mathcal{D}_c B_{j-1}] W_{j-1} N_{j-1}$ does not involve any $O(n)$ operation. Hence, this block can be computed quite cheaply. We give a brief overview of the cost of the RBiCG algorithm in Section 6; for a more detailed derivation, see [4].

- $\tilde{C}_{j-1}^* C = M_{j-1}^* \tilde{W}_{j-1}^* \begin{bmatrix} \tilde{C}_{j-2}^* C \\ \tilde{B}_{j-1}^* \mathcal{D}_c \end{bmatrix}.$

²For $p = 0$, no recycle space would be selected. This has never occurred in our experience. Indeed, discarding even one pair of vectors is rare.

The derivation of this block is similar to the previous block. As above, \tilde{C}_{j-2}^*C is available from the previous cycle, and \tilde{B}_{j-1} has been computed during the bi-Lanczos iteration.

- $$\begin{aligned} \tilde{C}_{j-1}^* \Upsilon_j &= \tilde{U}_{j-1}^* A \Upsilon_j = M_{j-1}^* \tilde{W}_{j-1}^* \begin{bmatrix} \tilde{U}_{j-2}^* \\ \tilde{V}_{j-1}^* \end{bmatrix} A \Upsilon_j = M_{j-1}^* \tilde{W}_{j-1}^* \begin{bmatrix} 0 \\ \tilde{V}_{j-1}^* A \Upsilon_j \end{bmatrix} \\ &= M_{j-1}^* \tilde{W}_{j-1}^* \begin{bmatrix} 0 \\ [\tilde{C} \tilde{C}^* A^* \tilde{V}_{j-1} + \tilde{\Upsilon}_{j-1} \tilde{\Gamma}_{j-1}]^* \Upsilon_j \end{bmatrix} \\ &= M_{j-1}^* \tilde{W}_{j-1}^* \begin{bmatrix} 0 \\ \tilde{\Gamma}_{j-1}^* \tilde{\Upsilon}_{j-1}^* \Upsilon_j \end{bmatrix}, \end{aligned}$$

where

$$\begin{aligned} \tilde{\Gamma}_{j-1}^* \tilde{\Upsilon}_{j-1}^* \Upsilon_j &= \tilde{\Gamma}_{j-1}^* \begin{bmatrix} \tilde{v}_{(j-2)s}^* \\ \tilde{v}_{(j-2)s+1}^* \\ \vdots \\ \tilde{v}_{(j-1)s}^* \\ \tilde{v}_{(j-1)s+1}^* \end{bmatrix} \begin{bmatrix} v_{(j-1)s} & v_{(j-1)s+1} & \cdots & v_{js} & v_{js+1} \end{bmatrix} \\ &= \tilde{\Gamma}_{j-1}^* \begin{bmatrix} 0 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & 0 \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \end{bmatrix}. \end{aligned}$$

- $$\tilde{\Upsilon}_j^* C_{j-1} = \begin{bmatrix} 0 & \tilde{\Upsilon}_j^* \Upsilon_{j-1} \Gamma_{j-1} \end{bmatrix} W_{j-1} N_{j-1}, \quad \text{where}$$

$$\tilde{\Upsilon}_j^* \Upsilon_{j-1} \Gamma_{j-1} = \begin{bmatrix} 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & \cdots & 0 & 1 \\ 0 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & 0 \end{bmatrix} \Gamma_{j-1}.$$

The derivation of this block is similar to the previous block.

- $$\tilde{C}^* U_{j-1} = \tilde{C}^* \begin{bmatrix} U_{j-2} & V_{j-1} \end{bmatrix} W_{j-1} N_{j-1} = \begin{bmatrix} \tilde{C}^* U_{j-2} & 0 \end{bmatrix} W_{j-1} N_{j-1},$$

where $\tilde{C}^* U_{j-2}$ is available from the previous cycle (such a block must be computed at most once per linear system).

- $$\begin{aligned} \tilde{C}_{j-1}^* U_{j-1} &= M_{j-1}^* \tilde{W}_{j-1}^* \begin{bmatrix} \tilde{U}_{j-2}^* \\ \tilde{V}_{j-1}^* \end{bmatrix} A \begin{bmatrix} U_{j-2} & V_{j-1} \end{bmatrix} W_{j-1} N_{j-1} \\ &= M_{j-1}^* \tilde{W}_{j-1}^* \begin{bmatrix} \tilde{C}_{j-2}^* U_{j-2} & \tilde{C}_{j-2}^* V_{j-1} \\ \tilde{V}_{j-1}^* C_{j-2} & \tilde{V}_{j-1}^* A V_{j-1} \end{bmatrix} W_{j-1} N_{j-1}, \end{aligned}$$

where $\tilde{C}_{j-2}^* U_{j-2}$ and $\tilde{C}_{j-2}^* V_{j-1}$ are submatrices of $\tilde{\Psi}_{j-1}^* \Phi_{j-1}$, $\tilde{V}_{j-1}^* C_{j-2}$ is a submatrix of $\tilde{\Upsilon}_{j-1}^* C_{j-2}$ and is available from $\tilde{\Psi}_{j-1}^* \Psi_{j-1}$, and $\tilde{V}_{j-1}^* A V_{j-1} = \tilde{T}_{j-1}$.

- $$\tilde{C}_{j-1}^* V_j \text{ is a submatrix of } \tilde{C}_{j-1}^* \Upsilon_j, \text{ and hence, is available from } \tilde{\Psi}_j^* \Psi_j.$$

Therefore, only $\tilde{\Upsilon}_j^* U_{j-1}$ needs to be computed explicitly.

5. Model Reduction. Consider a single-input/single-output (SISO) linear time-invariant (LTI) system represented as

$$G : \begin{cases} E \dot{x}(t) = Ax(t) + bv(t) \\ y(t) = c^*x(t), \end{cases} \quad \text{or} \quad G(s) = c^*(sE - A)^{-1}b \quad (5.1)$$

where $E, A \in \mathbb{R}^{n \times n}$ and $b, c \in \mathbb{R}^n$. The time-dependent functions $v(t), y(t) : \mathbb{R} \rightarrow \mathbb{R}$ are the input and output of $G(s)$, respectively, and $x(t) : \mathbb{R} \rightarrow \mathbb{R}^n$ is the associated state. In (5.1), $G(s)$ is the transfer function of the system: Let $V(s)$ and $Y(s)$ denote the Laplace transforms of $v(t)$ and $y(t)$, respectively. Then, the transfer function $G(s)$ satisfies $Y(s) = G(s)V(s)$. By a common abuse of notation, we denote both the underlying dynamical system and its transfer function with G . The dimension of the underlying state-space, n , is called the dimension of G . Systems of the form (5.1) with extremely large state-space dimension n arise in many applications; see [6] and [33] for a collection of such examples. Simulations in such large scale settings lead to overwhelming demands on computational resources. This is the main motivation for model reduction. The goal is to produce a surrogate model of much smaller dimension which provides a high-fidelity approximation of the input-output behavior of the original model G . Let $r \ll n$ denote the order of the reduced-model. The reduced-model is represented, similar to (5.1), as

$$G_r(s) : \begin{cases} E_r \dot{x}_r(t) = A_r x_r(t) + b_r v(t) \\ y_r(t) = c_r^* x_r(t), \end{cases} \quad \text{or} \quad G_r(s) = c_r^*(sE_r - A_r)^{-1}b_r \quad (5.2)$$

where $E_r, A_r \in \mathbb{R}^{r \times r}$ and $b_r, c_r \in \mathbb{R}^r$. In this setting, the common approach is to construct reduced order models via a Petrov-Galerkin projection. This amounts to choosing two r -dimensional subspaces \mathcal{V}_r and \mathcal{W}_r and matrices $V_r \in \mathbb{R}^{n \times r}$ and $W_r \in \mathbb{R}^{n \times r}$ such that $\mathcal{V}_r = \text{Range}(V_r)$ and $\mathcal{W}_r = \text{Range}(W_r)$. Then, we approximate the full-order state $x(t)$ as $x(t) \approx V_r x_r(t)$ and enforce the Petrov-Galerkin condition,

$$W_r^* (E V_r \dot{x}_r(t) - A V_r x_r(t) - b v(t)) = 0, \quad y_r(t) = c^* V_r x_r(t),$$

leading to a reduced-order model as in (5.2) with

$$E_r = W_r^* E V_r, \quad A_r = W_r^* A V_r, \quad b_r = W_r^* b, \quad \text{and} \quad c_r = V_r^* c. \quad (5.3)$$

As (5.3) illustrates, the quality of the reduced model depends solely on the selection of the two subspaces \mathcal{V}_r and \mathcal{W}_r . In this paper, we will choose \mathcal{V}_r and \mathcal{W}_r to enforce interpolation. For other selections of \mathcal{V}_r and \mathcal{W}_r , we refer the reader to [6].

5.1. Interpolatory Model Reduction. For a given full-order model $G(s)$, the goal of interpolatory model reduction is to construct a reduced-order model $G_r(s)$ via rational interpolation. Here, we focus on Hermite interpolation. Given the full-order model (5.1) and a collection of interpolation points (also called shifts) $\sigma_i \in \mathbb{C}$, for $i = 1, \dots, r$, we must construct a reduced-order system by projection as in (5.3) such that $G_r(s)$ interpolates $G(s)$ and its first derivative at selected interpolation points, i.e.,

$$G(\sigma_i) = G_r(\sigma_i) \quad \text{and} \quad G'(\sigma_i) = G'_r(\sigma_i) \quad \text{for} \quad i = 1, \dots, r.$$

Rational interpolation by projection was first proposed in [17, 55, 56]. How to obtain the required projection was derived in [27] using the rational Krylov method [40]. For

the special case of Hermite rational interpolation, the solution of the interpolatory model reduction problem is given in Theorem 5.1. For the more general case, we refer the reader to [27] and the recent survey [7].

THEOREM 5.1. *Given $G(s) = c^*(sE - A)^{-1}b$ and r distinct points $\sigma_1, \dots, \sigma_r \in \mathbb{C}$, let*

$$V_r = [(\sigma_1 E - A)^{-1}b \dots (\sigma_r E - A)^{-1}b], \quad W_r^* = \begin{bmatrix} c^*(\sigma_1 E - A)^{-1} \\ \vdots \\ c^*(\sigma_r E - A)^{-1} \end{bmatrix}. \quad (5.4)$$

Using (5.3), define the reduced-order model $G_r(s) = c_r^(sE_r - A_r)^{-1}b_r$. Then $G(\sigma_i) = G_r(\sigma_i)$ and $G'(\sigma_i) = G_r'(\sigma_i)$, for $i = 1, \dots, r$, provided that $\sigma_i E - A$ and $\sigma_i E_r - A_r$ are invertible for $i = 1, \dots, r$.*

Theorem 5.1 shows how to solve the interpolatory model reduction problem via projection for given shifts. However, it does not provide a strategy for choosing good/optimal interpolation points. Recently, this issue has been resolved for the special case of optimality in the \mathcal{H}_2 norm [29]. The \mathcal{H}_2 norm of the dynamical system $G(s)$ is defined as

$$\|G\|_{\mathcal{H}_2} = \left(\frac{1}{2\pi} \int_{-\infty}^{\infty} |G(ja)|^2 da \right)^{1/2}.$$

The \mathcal{H}_2 norm of G is the $2 - \infty$ induced norm of the underlying convolution operator. Then, for any $v \in L^2(\mathbb{R}^+)$, $\|y - y_r\|_{L^\infty} \leq \|G - G_r\|_{\mathcal{H}_2} \|v\|_{L^2}$. To ensure that the output error $y - y_r$ is small in $L^\infty(\mathbb{R}^+)$ uniformly over all inputs v , say, with $\|v\|_{L^2} \leq 1$, we seek a reduced system G_r that makes $\|G - G_r\|_{\mathcal{H}_2}$ small. This leads to the *optimal \mathcal{H}_2 model reduction problem*: Given $G(s)$, and a reduced order $r < n$, find $G_r(s)$ that solves

$$\|G - G_r\|_{\mathcal{H}_2} = \min_{\dim(\hat{G}_r)=r} \|G - \hat{G}_r\|_{\mathcal{H}_2}. \quad (5.5)$$

This problem has been studied extensively [35, 54, 29, 46, 52, 28, 10, 11]. It is a non-convex optimization problem, which makes finding the global minimum, at best, a hard task. Hence, the common approach is to construct reduced-order models that satisfy, for an interpolatory model reduction framework, the following first-order necessary conditions.

THEOREM 5.2. ([35, 29]) *Given $G(s)$, let $G_r(s) = c_r^*(sE_r - A_r)^{-1}b_r$ be an \mathcal{H}_2 -optimal reduced order model of order r , and let $\hat{\lambda}_1, \dots, \hat{\lambda}_r$ denote the poles of $G(s)$. Then*

$$G(-\hat{\lambda}_i) = G_r(-\hat{\lambda}_i) \quad \text{and} \quad G'(-\hat{\lambda}_i) = G_r'(-\hat{\lambda}_i) \quad \text{for } i = 1, \dots, r. \quad (5.6)$$

So, the \mathcal{H}_2 optimal approximant $G_r(s)$ is a Hermite interpolant to $G(s)$ at the mirror image of its poles. These poles, the optimal interpolation points, are not known *a priori*. Hence, the iterative rational Krylov algorithm (IRKA) [29], starting from an initial selection of interpolation points, iteratively corrects the interpolation points until (5.6) is satisfied. Algorithm 3 outlines IRKA; for details, see [29].

5.2. Advantages of Approximating Solutions using a Petrov-Galerkin Framework in Interpolatory Model Reduction. The main cost in IRKA is solving multiple linear systems to compute V_r and W_r . If the dimension of the state-space,

Algorithm 3. *IRKA* ([29])

1. Make an initial shift selection σ_i for $i = 1, \dots, r$,
2. $V_r = [(\sigma_1 E - A)^{-1}b, \dots, (\sigma_r E - A)^{-1}b]$,
3. $W_r = [(\sigma_1 E - A)^{-*}c, \dots, (\sigma_r E - A)^{-*}c]$,
4. while (not converged)
 - ◇ $A_r = W_r^* A V_r$, $E_r = W_r^* A V_r$,
 - ◇ $\sigma_i \leftarrow -\lambda_i(A_r, E_r)$ for $i = 1, \dots, r$,
 - ◇ $V_r = [(\sigma_1 E - A)^{-1}b, \dots, (\sigma_r E - A)^{-1}b]$,
 - ◇ $W_r = [(\sigma_1 E - A)^{-*}c, \dots, (\sigma_r E - A)^{-*}c]$,
5. $A_r = W_r^* A V_r$, $E_r = W_r^* A V_r$, $b_r = W_r^* b$, $c_r = V_r^* c$.

n , is large, these systems are generally solved only approximately by an iterative solver. In this context, it is important to assess the accuracy of the computed reduced order model; that is, given the shifts, how accurately the Hermite interpolation problem is solved. This question was studied extensively in [12]. One of the major results, outlined below for our particular case, provides the main motivation for solving the linear systems associated with the corresponding columns of V_r and W_r as pairs of dual linear systems (in the terminology of Section 1) using BiCG or RBiCG.

Let \hat{v}_j and \hat{w}_j , for $j = 1, \dots, r$, denote the approximate solutions of $(\sigma_j E - A)v_j = b$ and $(\sigma_j E - A)^* w_j = c$, respectively, with residuals $\eta_j = (\sigma_j E - A)\hat{v}_j - b$ and $\xi_j = (\sigma_j E - A)^* \hat{w}_j - c$. Furthermore, let \hat{v}_j , \hat{w}_j , η_j , and ξ_j satisfy the Petrov-Galerkin condition that there exist spaces \mathcal{P} and \mathcal{Q} such that $\hat{v}_j \in \mathcal{P}$, $\hat{w}_j \in \mathcal{Q}$, $\eta_j \perp \mathcal{Q}$, and $\xi_j \perp \mathcal{P}$. Define the approximate solution matrices (\hat{V}_r and \hat{W}_r), the residual matrices (R_b and R_c), and the rank-2r matrix (F_{2r}) as follows:

$$\begin{aligned} \hat{V}_r &= [\hat{v}_1 \ \hat{v}_2 \ \dots \ \hat{v}_r], & \hat{W}_r &= [\hat{w}_1 \ \hat{w}_2 \ \dots \ \hat{w}_r], \\ R_b &= [\eta_1 \ \eta_2 \ \dots \ \eta_r], & R_c &= [\xi_1 \ \xi_2 \ \dots \ \xi_r], \\ F_{2r} &= R_b(\hat{W}_r^* \hat{V}_r)^{-1} \hat{W}_r^* + \hat{V}_r^* (\hat{W}_r^* \hat{V}_r)^{-1} R_c. \end{aligned}$$

Also, define the inexact reduced-order order quantities

$$\hat{A}_r = \hat{W}_r^* A \hat{V}_r, \quad \hat{E}_r = \hat{W}_r^* E \hat{V}_r, \quad \hat{b}_r = \hat{W}_r^* b, \quad \text{and} \quad \hat{c}_r = \hat{V}_r^* c.$$

Then, the computed reduced-order model $\hat{G}_r(s) = \hat{c}_r^* (s \hat{E}_r - \hat{A}_r)^{-1} \hat{b}_r$ *exactly* interpolates the perturbed full-order model $\hat{G}(s) = c^* (sE - (A + F_{2r}))^{-1} b$, i.e.,

$$\hat{G}(\sigma_i) = \hat{G}_r(\sigma_i) \quad \text{and} \quad \hat{G}'(\sigma_i) = \hat{G}'_r(\sigma_i), \quad \text{for } i = 1, \dots, r.$$

Hence, iteratively solving the linear systems while satisfying the Petrov-Galerkin condition above yields a backward error for the interpolatory model reduction that is bounded by $\|F_{2r}\|$, which is governed by the norms of the residuals. The latter are easily controlled in the iterative solver. For details, we refer the reader to [12].

The easiest way to satisfy the Petrov-Galerkin condition above is by solving the dual pairs of linear systems using BiCG. Hence, BiCG is particularly suitable for solving the linear systems in IRKA (even for symmetric positive definite matrices). However, as IRKA leads to a sequence of dual linear systems, the RBiCG algorithm can be used to reduce the total run time for solving all linear systems. Moreover, if we solve the dual pairs of linear systems arising in IRKA by RBiCG, the Petrov-Galerkin condition is still satisfied. Hence, the resulting reduced-order model will be an optimal \mathcal{H}_2 approximation to a nearby full-order model.

5.3. IRKA using RBiCG. IRKA usually converges rather fast [29]. Hence, after one or a few initial steps, the interpolation points from one step of IRKA to the next do not change substantially. Moreover, for many cases, the change of the (appropriately ordered) $\{\sigma_i\}$ from one column of V_r (and W_r) to the next is also modest. Therefore, IRKA is expected to gain significantly from recycling.

For the special case of $E = I$ in (5.1), alternative solution approaches might be advantageous, as one can solve the linear systems for multiple shifts at once [21, 25, 31, 49]. Combining these strategies with a Petrov-Galerkin framework does not seem complicated. Effective strategies for Krylov subspace recycling for solving systems of the type, $(\sigma_i I - A)v_i = b$, for multiple shifts at once, as well as for multiple right hand sides, was discussed in [32]. For most model reduction applications, however, $E \neq I$.

There are three strategies for recycling Krylov subspaces in IRKA. For the first strategy, consider two consecutive steps of IRKA, say step m and $m + 1$ (iterations m and $m + 1$ of the while loop in Algorithm 3), with shifts $\sigma_i^{(m)}$, for $i = 1, \dots, r$ and

$$\begin{aligned} V_r^{(m)} &= [(\sigma_1^{(m)} E - A)^{-1}b, \dots, (\sigma_r^{(m)} E - A)^{-1}b], \\ W_r^{(m)} &= [(\sigma_1^{(m)} E - A)^{-*}c, \dots, (\sigma_r^{(m)} E - A)^{-*}c]. \end{aligned} \quad (5.7)$$

at step m and with shifts $\sigma_i^{(m+1)}$, for $i = 1, \dots, r$ and

$$\begin{aligned} V_r^{(m+1)} &= [(\sigma_1^{(m+1)} E - A)^{-1}b, \dots, (\sigma_r^{(m+1)} E - A)^{-1}b], \\ W_r^{(m+1)} &= [(\sigma_1^{(m+1)} E - A)^{-*}c, \dots, (\sigma_r^{(m+1)} E - A)^{-*}c] \end{aligned}$$

at step $m + 1$ of IRKA. One can recycle Krylov subspaces from the i^{th} column of $V_r^{(m)}$ and $W_r^{(m)}$ to the i^{th} column of $V_r^{(m+1)}$ and $W_r^{(m+1)}$. That is, from solving the pair of linear systems

$$(\sigma_i^{(m)} E - A)v_i^{(m)} = b, \quad (\sigma_i^{(m)} E - A)^* w_i^{(m)} = c,$$

to solving the pair of linear systems

$$(\sigma_i^{(m+1)} E - A)v_i^{(m+1)} = b, \quad (\sigma_i^{(m+1)} E - A)^* w_i^{(m+1)} = c,$$

where $i = 1, 2, \dots, r$. This strategy for recycling strategy is useful when the change in a shift from one IRKA step to the next is small.

For the second strategy, consider a single IRKA step. One can recycle selected Krylov subspaces from solving for one pair of columns of V_r and W_r to the next pair of columns across all the columns of the matrices V_r and W_r . In the third strategy, the first two recycling strategies are combined. We describe one such combination. Consider solving the system $(\sigma_i^{(m+1)} E - A)v_i^{(m+1)} = b$ and its dual system. From a set of previously generated recycle spaces (distinguished by their shifts), one can pick the recycle space from the system with the smallest relative change in σ (and less than a relative tolerance). This would ensure that the linear system from which the recycle space has been generated is close to the current one. A natural pool from which to pick the σ defining the recycle space would be $\sigma_1^{(m)}, \dots, \sigma_r^{(m)}, \sigma_1^{(m+1)}, \dots, \sigma_{i-1}^{(m+1)}$. The second and third recycling strategies are useful when the shifts at an IRKA step are clustered.

For the experiments in this paper, r is small, and so the shifts at any particular IRKA step are spread far apart. Hence, we follow the first strategy. That is, for every shift, we recycle Krylov subspaces from one IRKA step to the next. In general,

the linear systems corresponding to the relatively large shifts converge fast, and so recycling Krylov subspaces is not useful for them. Therefore, we carry out recycling only for selected, small shifts. We give more details in Section 6.2.

5.4. Previous Work in Recycling for Model Reduction. Recycling for interpolatory model reduction in the Galerkin setting, i.e., with $W_r = V_r$, has been considered in [14] and [19]. In this setting, there are no dual systems to solve, and therefore approaches based on GCR [18] and GMRES [42] are considered, respectively, for a sequence of (single) linear systems, as opposed to our approach based on BiCG for a sequence of dual linear systems. Also in other respects, the approach for improving the linear solver and the model reduction context are quite different from here. In [14], the focus is on efficiently solving linear systems with a fixed coefficient matrix and multiple right hand sides ($Ax^{(j)} = b^{(j)}$), recycling descent vectors (in GCR). Furthermore, the authors target model reduction with a single interpolation point but interpolating higher derivatives.

6. Results. We first give a brief overview of the overhead in RBiCG. We focus on components with at least $O(n)$ cost, where n is the dimension of the linear system. Furthermore, k is the number of basis vectors in the primal (or dual) recycle space, and s is the number of iterations in a cycle. For every iteration, there is an extra cost of $(8k + 2)n$ flops, mostly from orthogonalizations. At the end of each cycle, there is an extra cost of $(14k^2 + 6ks + 16k + 4)n$ flops, mostly from setting up the generalized eigenvalue problem and computing biorthogonal C_j and \tilde{C}_j . Once per linear system, there is an extra cost of $(10k^2 + 28k + 14)n$ flops, mostly from computing biorthogonal C and \tilde{C} . A more detailed discussion of the overhead is given in [4]. Note that s and k are much smaller than n . For recycling to be beneficial, the savings in iterations should be sufficient to make up for the overhead. Further in this section, we show that the reduction in the number of iterations for (a pair of) linear systems may be as high as 70%. For our model reduction test problem, we show that computing a reduced model without recycling takes about 50% more time than with recycling.

We test RBiCG on a convection-diffusion problem and on IRKA for interpolatory model reduction. All experiments are done using Matlab.

6.1. Convection-Diffusion. To analyze RBiCG, we use the linear system obtained by finite difference discretization of the partial differential equation

$$-(\mathcal{A}\vartheta_x)_x - (\mathcal{A}\vartheta_y)_y + \mathcal{B}(x, y)\vartheta_x = \mathcal{F},$$

with \mathcal{A} as shown in Figure 6.1 (a), $\mathcal{B}(x, y) = 2e^{2(x^2+y^2)}$, and $\mathcal{F} = 0$ everywhere except in a small square in the center where $\mathcal{F} = 100$ [50]; see Figure 6.1(a). The domain is $(0, 1) \times (0, 1)$ with Dirichlet boundary conditions

$$\vartheta(0, y) = \vartheta(1, y) = \vartheta(x, 0) = 1, \text{ and } \vartheta(x, 1) = 0.$$

We use the second order central difference scheme with a mesh width of $h = 1/64$, giving a nonsymmetric linear system of 3969 unknowns. The convergence is similar for a problem that is four times larger. To enable further analysis, we give results for this smaller system size. The primary system right-hand side comes from the PDE. We take the vector of all zeros as the dual system right-hand side. In this case, we are concerned only about the primary system.

To analyze RBiCG, we solve the (same) dual linear systems four times. The recycle space generated during the first run is used for solving the same dual systems

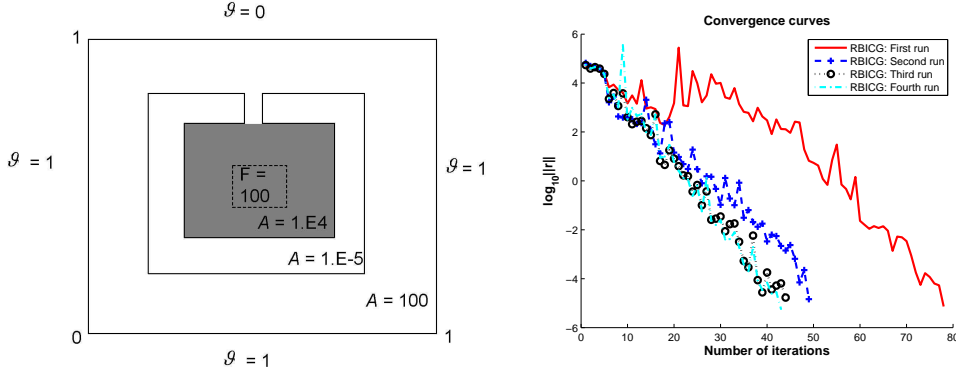


FIG. 6.1. RBiCG for a convection-diffusion problem. (a) The coefficients of the PDE. (b) Convergence for preconditioned RBiCG with $s = 40$ and $k = 10$ for the primary system solved four times to analyze convergence improvement as the recycle space improves.

a second time, further improving the recycle space, and so on. This is a useful approach for analyzing how well Krylov subspace recycling works, as it excludes the effects of changing matrices and of right-hand sides having different expansions in the eigenvector basis [38]. Hence, it provides an indication for reasonable sequences of systems how fast the recycle spaces converge and how much recycling approximate invariant subspaces is likely to improve convergence. For this experiment, we take $s = 40$ and $k = 10$. These are chosen based on experience with other recycling algorithms [38, 53]. The relative tolerance for RBiCG is taken as 10^{-8} . The initial guess (for both systems) is a vector of all ones. The linear systems are split-preconditioned by a Crout version of the ILUT preconditioner with a drop tolerance of 0.05 [41]. The generated recycle spaces pertain to the preconditioned linear systems.

Figure 6.1(b) shows the convergence improvement of RBiCG, as it solves the primary system multiple times. For the second run, the reduction in iterations is around 35%. The convergence improves further with each run. Next, we present a brief analysis of the generated recycle spaces. In Table 6.1, we give the cosines of the principal angles between the primary (dual) recycle space and the right (left) invariant subspace associated with the eight eigenvalues of smallest magnitude. As for the recycle spaces, the invariant subspaces are computed for the preconditioned operator. As the recycle space improves, the principal angles tend to zero, and so the cosines tend to one. The table shows that with only a few runs, RBiCG accurately approximates increasingly larger subspaces of the invariant subspace. As a result, we see faster convergence for every new run.

6.2. Model Reduction. Our test dynamical system is a semi-discretized heat transfer problem for determining the optimal cooling of steel profiles [39, 13, 43]. We will refer to this model as the *rail model* [39]. The rail model has seven inputs and six outputs. Since we focus on SISO systems in this paper, we choose a SISO subsystem corresponding to the second input and sixth output. The rail model is available with 1357, 5177, 20209, and 79841 unknowns, depending on the mesh size.

As convergence tolerance for IRKA we use a relative change in the shifts of less than 10^{-6} . The matrices A and E of (5.1) are symmetric negative definite and symmetric positive definite (SPD), respectively. Since our shifts are real and positive at every IRKA step, $(\sigma_i^{(m)} E - A)$ is always SPD. Nevertheless, RBiCG is advantageous

Primary System			Dual System		
Start of Run 2	Start of Run 3	Start of Run 4	Start of Run 2	Start of Run 3	Start of Run 4
1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
0.9896	1.0000	1.0000	0.9950	1.0000	1.0000
0.3832	1.0000	1.0000	0.9884	1.0000	1.0000
0.1452	0.9983	1.0000	0.7864	0.9844	1.0000
0.0988	0.9437	0.9970	0.6070	0.9206	0.8141
0.0300	0.1869	0.9567	0.4749	0.4118	0.4721

TABLE 6.1

Convergence of the recycle space for the convection-diffusion problem as measured by the cosines of the principal angles between the primary (dual) recycle space and the right (left) invariant subspace associated with the eight eigenvalues of smallest magnitude.

here because of the backward error formulation discussed in Section 5.2. We carry out two sets of experiments that differ in the dimension, r , of the reduced models. We also vary the frequency of computing a recycle space, as a recycle space can be effective for multiple consecutive systems [38, 32] and updating it may be expensive.

We implement the first recycling strategy from Section 5.3 for a few selected shifts. As for the convection-diffusion example, the recycling parameters s and k are chosen based on experience with other recycling algorithms [38, 53]. If a pair of linear systems converges in fewer than s iterations, the recycle space is not updated, and we use the previous recycle space for the next pair of systems in the sequence. The relative convergence tolerance for the iterative solves and the tolerance for constructing \tilde{C}_j and C_j in Section 4.3 are taken as 10^{-6} . The linear systems are split-preconditioned with an incomplete LU preconditioner with threshold and pivoting (ILUTP) [41]. The drop tolerance varies per problem to avoid ill-conditioning; see Figures 6.2 – 6.5. The initial guess of the preconditioned system is the solution vector from the previous preconditioned system in the sequence. For the first IRKA step, we take a vector of all zeros as the initial guess. In general, a better initial guess may be based on knowledge of the system and aim to avoid orthogonal initial residuals (Algorithm 1 Step 2; Algorithm 2 Step 3).

For the first set of experiments, we reduce the models to $r = 6$ degrees of freedom, with 1.00×10^{-5} , 1.38×10^{-4} , 1.91×10^{-3} , 2.63×10^{-2} , 3.63×10^{-1} , and 5.01 as initial shifts.

We compute a recycle space at every IRKA step. The results for the primary systems at a particular IRKA step (given in the caption) are given in Figures 6.2 – 6.5. The graphs for the other IRKA steps are similar, as are the graphs for the dual systems. We carry out recycling for the smallest two shifts. Each figure has two solid curves for the linear systems solved without recycling and two dashed-dotted curves for those solved with recycling. It is evident that recycling significantly reduces the number of iterations. The savings in iterations are as high as 70% per system. As discussed in Section 5.3, convergence for the remaining four (larger) shifts is rapid, so recycling Krylov subspaces is not useful for these.

Next, we analyze the recycle space generated during the first two IRKA steps for the order 5177 rail model corresponding to the smallest shift. In Table 6.2, we give the cosines of principal angles between the recycle space and the invariant subspace




FIG. 6.2. Convergence of preconditioned $RBiCG$ at the 3rd IRKA step for the $n = 1357$ rail model, with $s = 40$, $k = 10$, and the preconditioner drop tolerance is 0.1.




FIG. 6.3. Convergence of preconditioned $RBiCG$ at the 2nd IRKA step for the $n = 5177$ rail model, with $s = 50$, $k = 10$, and the preconditioner drop tolerance is 0.05.

spanned by eight eigenvectors associated with the eigenvalues of smallest magnitude. As for the recycle space, the invariant subspace is computed for the preconditioned operator. For the primary system, we use the right invariant subspace. For the dual system, we use the left invariant subspace. As the recycle space improves, the principal angles tend to zero, and so the cosines tend to one. Consider the results for the primary system. At the first IRKA step and the end of the first cycle, we see that the recycle space captures a subspace of dimension four of the invariant subspace. The recycle space gets more accurate at the end of the second cycle and captures a subspace of dimension seven. For the second IRKA step, we have a new shift, and so the matrix changes. Therefore, at the start of the first cycle, we see a slight deterioration of the recycle space (almost negligible). This recycle space leads to the

Primary System				Dual System			
IRKA Step 1 $\sigma_1 = 1.000 \times 10^{-5}$		IRKA Step 2 $\sigma_1 = 1.834 \times 10^{-5}$		IRKA Step 1 $\sigma_1 = 1.000 \times 10^{-5}$		IRKA Step 2 $\sigma_1 = 1.834 \times 10^{-5}$	
End of Cycle 1	End of Cycle 2	Start of Cycle 1	End of Cycle 1	End of Cycle 1	End of Cycle 2	Start of Cycle 1	End of Cycle 1
1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
1.0000	1.0000	1.0000	1.0000	0.9997	1.0000	1.0000	1.0000
0.9987	1.0000	1.0000	1.0000	0.9765	1.0000	1.0000	1.0000
0.9321	1.0000	1.0000	1.0000	0.4936	1.0000	1.0000	1.0000
0.2257	1.0000	0.9998	0.9999	0.0844	0.9995	0.9997	0.9998
0.0260	0.9997	0.9996	0.9997	0.0231	0.9945	0.9945	0.9989
0.0072	0.7813	0.7799	0.9932	0.0068	0.3439	0.3423	0.9876

TABLE 6.2

Convergence of the recycle space for the sequence of linear systems corresponding to the 5177 rail model and the smallest shift, as measured by the cosines of the principal angles between the primary (dual) recycle space and the right (left) invariant subspace associated with the eight eigenvalues of smallest magnitude. The third column corresponds to the dashed convergence curve in Figure 6.3.




FIG. 6.4. Convergence of preconditioned RBiCG at the 2nd IRKA step for the $n = 20209$ rail model, with $s = 40$, $k = 20$, and the preconditioner drop tolerance is 0.01.

dashed curve in Figure 6.3. By the end of the first cycle (at the second IRKA step), all eight eigenvectors are captured. The results for the dual system recycle space are similar.

For the second set of experiments, we reduce the models to $r = 3$ degrees of freedom, using as initial shifts, 1.00×10^{-5} , 7.08×10^{-3} , and 5.01. We compute the recycle space at every fifth IRKA step. The results are given in Table 6.3. We implement recycling for the smallest shift only. The linear systems corresponding to the two (larger) shifts converge fast, so recycling Krylov subspaces is not useful for these. *Total iteration count* refers to the sum of iteration counts for solving linear systems over all shifts and all IRKA steps. *Total time* is the time in seconds required by IRKA




FIG. 6.5. Convergence of preconditioned RBiCG at the 3rd IRKA step for the $n = 79841$ rail model, with $s = 50$, $k = 20$, and the preconditioner drop tolerance is 0.005.

to converge to the ideal shifts. This includes the time for all IRKA computations as well as all linear solves (BiCG or RBiCG, as the case may be). The table illustrates that computing the reduced model without recycling takes about 50% more time than with recycling. Obviously, the improvement for just the pairs of linear systems where recycling is actually used is substantially larger.

Size	s	k	Drop tol	IRKA steps	Total iteration count			Total time (s)		
					BiCG	RBiCG	Ratio	BiCG	RBiCG	Ratio
20209	40	20	0.01	31	3032	1434	2.11	73.82	54.28	1.36
79841	50	20	0.005	44	6324	2547	2.48	742.83	505.09	1.47

TABLE 6.3

The total number of iterations and computation time over all IRKA iterations with BiCG and with RBiCG for the linear systems with the smallest shift. Total time includes the time for all computations to compute the reduced model.

7. Conclusion. We focus on efficiently solving sequences of dual linear systems. For several classes of problems, such as the linear systems arising in interpolatory model reduction, or bilinear forms arising in Quantum Monte Carlo methods, the BiCG algorithm has advantages over methods like GMRES that would solve the primary and the dual system separately. For sequences of dual linear systems arising in such problems, it is advantageous to use Krylov subspace recycling for the BiCG algorithm, and for this purpose we propose the RBiCG algorithm. The derivation of RBiCG also provides the foundation for recycling variants of other popular bi-Lanczos based methods, like CGS, BiCGSTAB, QMR, and TFQMR [3].

We have demonstrated the usefulness of RBiCG for interpolatory model reduction using IRKA, an application that may be an important niche for this solver. In addition, we have analyzed and demonstrated the effectiveness of RBiCG for non-

symmetric linear systems arising from convection-diffusion problems. This suggests that the RBiCG method may be useful in other areas where solving dual systems in a Petrov-Galerkin sense brings special advantages.

In future work, we plan to extend the use of RBiCG to model reduction for MIMO dynamical systems in a tangential interpolation framework where the right-hand sides are not constant as in the SISO case. In addition, we will investigate the use of RBiCG for evaluating bilinear forms arising in QMC algorithms. Our current results for this look promising.

Acknowledgments. We thank the anonymous reviewers for their careful and helpful suggestions, which greatly helped us to improve this paper.

REFERENCES

- [1] A. M. Abdel-Rehim, R. B. Morgan, and W. Wilcox. Deflated BiCGStab for linear equations in QCD problems. Technical report, arXiv:0710.1988, 2007.
- [2] A. M. Abdel-Rehim, A. Stathopoulos, and K. Orginos. Extending the eigCG algorithm to non-symmetric Lanczos for linear systems with multiple right-hand sides. Technical Report WM-CS-2009-06, College of William and Mary, 2009.
- [3] K. Ahuja. Recycling bi-Lanczos algorithms: BiCG, CGS, and BiCGSTAB. Master's thesis, Department of Mathematics, Virginia Tech, 2009. Advised by E. de Sturler. Available from <http://scholar.lib.vt.edu/theses/available/etd-08252009-161256/>.
- [4] K. Ahuja. *Recycling Krylov Subspaces and Preconditioners*. PhD thesis, Virginia Tech, October 2011. Advisor: Eric de Sturler.
- [5] K. Ahuja, B. K. Clark, E. de Sturler, D. M. Ceperley, and J. Kim. Improved scaling for quantum Monte Carlo on insulators. *SIAM Journal on Scientific Computing*, 33(4):1837–1859, 2011. Earlier versions available from <http://arxiv.org/abs/1008.5113>.
- [6] A. C. Antoulas. *Approximation of Large-Scale Dynamical Systems (Advances in Design and Control)*. SIAM, Philadelphia, PA, USA, 2005.
- [7] A. C. Antoulas, C. A. Beattie, and S. Gugercin. Interpolatory model reduction of large-scale dynamical systems. In J. Mohammadpour and K. Grigoriadis, editors, *Efficient Modeling and Control of Large-Scale Systems*. Springer-Verlag, 2010.
- [8] R. E. Bank and T. F. Chan. An analysis of the composite step biconjugate gradient method. *Numer. Math.*, 66:295–319, 1993.
- [9] C. A. Beattie. Harmonic Ritz and Lehmann bounds. *Electronic Transactions on Numerical Analysis*, 7:18–39, 1998.
- [10] C. A. Beattie and S. Gugercin. Krylov-based minimization for optimal \mathcal{H}_2 model reduction. In *Proceedings of the 46th IEEE Conference on Decision and Control*, pages 4385–4390, 2007.
- [11] C. A. Beattie and S. Gugercin. A trust region method for optimal \mathcal{H}_2 model reduction. In *Proceedings of 48th IEEE Conference on Decision & Control and 28th Chinese Control Conference*, pages 5370–5375, 2009.
- [12] C. A. Beattie, S. Gugercin, and S. Wyatt. Inexact solves in interpolatory model reduction. *Linear Algebra and its Applications (accepted)*, 2010. Available as arXiv:1007.5213v2.
- [13] P. Benner. Solving large-scale control problems. *IEEE Control Systems Magazine*, 24(1):44–59, 2004.
- [14] P. Benner and L. Feng. On recycling Krylov subspaces for solving linear systems with successive right-hand sides with applications in model reduction. In P. Benner, M. Hinze, and E. J. W. ter Maten, editors, *Model Reduction for Circuit Simulation*, volume 74 of *Lecture Notes in Electrical Engineering*. Springer-Verlag, 2011 (in press).
- [15] E. de Sturler. BiCG explained. In *Householder Symposium XIV, Proceedings of the Householder International Symposium in Numerical Algebra, Chateau Whistler, Whistler, BC, Canada, June 13-19, 1999*, 1999.
- [16] E. de Sturler. Truncation strategies for optimal Krylov subspace methods. *SIAM Journal on Numerical Analysis*, 36(3):864–889, 1999.
- [17] C. De Villemagne and R. E. Skelton. Model reductions using a projection formulation. *International Journal of Control*, 46(6):2141–2169, 1987.
- [18] S. C. Eisenstat, H. C. Elman, and M. H. Schultz. Variational iterative methods for nonsymmetric systems of linear equations. *SIAM Journal on Numerical Analysis*, 20(2):345–357,

- 1983.
- [19] L. Feng, P. Benner, and J. Korvink. Parametric model order reduction accelerated by subspace recycling. In *Proceedings of 48th IEEE Conference on Decision & Control and 28th Chinese Control Conference*, pages 4328–4333, 2009.
- [20] R. Fletcher. Conjugate gradient methods for indefinite systems. *Lecture Notes in Mathematics, Springer Berlin-Heidelberg*, 506:73–89, 1976.
- [21] R. W. Freund. Solution of shifted linear systems by quasi-minimal residual iterations. In L. Reichel, A. Ruttan, and R. S. Varga, editors, *Numerical Linear Algebra*, pages 101–121. W. de Gruyter, 1993.
- [22] R. W. Freund. A transpose-free quasi-minimal residual algorithm for non-Hermitian linear systems. *SIAM Journal on Scientific Computing*, 14(2):470–482, 1993.
- [23] R. W. Freund, M. H. Gutknecht, and N. M. Nachtigal. An implementation of the look-ahead Lanczos algorithm for non-Hermitian matrices. *SIAM Journal on Scientific Computing*, 14(1):137–158, 1993.
- [24] R. W. Freund and N. M. Nachtigal. QMR: a quasi-minimal residual method for non-Hermitian linear systems. *Numer. Math.*, 60:315–339, 1991.
- [25] A. Frommer. BiCGStab(ℓ) for families of shifted linear systems. *Computing*, 70:87–109, 2003.
- [26] A. Greenbaum. *Iterative Methods for Solving Linear Systems*. SIAM, Philadelphia, PA, USA, 1997.
- [27] E. Grimme. *Krylov projection methods for model reduction*. PhD thesis, University of Illinois, 1997.
- [28] S. Gugercin. An iterative rational Krylov algorithm (IRKA) for optimal \mathcal{H}_2 model reduction. In *Householder Symposium XVI*, Seven Springs Mountain Resort, PA, USA, May 2005.
- [29] S. Gugercin, A. C. Antoulas, and C. A. Beattie. \mathcal{H}_2 model reduction for large-scale linear dynamical systems. *SIAM Journal on Matrix Analysis and Applications*, 30(2):609–638, 2008.
- [30] M. H. Gutknecht. Lanczos-type solvers for nonsymmetric linear systems of equations. *Acta Numerica*, 6:271–397, 1997.
- [31] B. Jegerlehner. Krylov space solvers for shifted linear systems. *Hep-lat/9612014*, 1996.
- [32] M. E. Kilmer and E. de Sturler. Recycling subspace information for diffuse optical tomography. *SIAM Journal on Scientific Computing*, 27(6):2140–2166, 2006.
- [33] J. Korvink and E. Rudnyi. Oberwolfach benchmark collection. In P. Benner, V. Mehrmann, and D. C. Sorensen, editors, *Dimension Reduction of Large-Scale Systems*, volume 45 of *Lecture Notes in Computational Science and Engineering*, pages 311–315. Springer-Verlag, Berlin/Heidelberg, Germany, 2005.
- [34] C. Lanczos. Solution of systems of linear equations by minimized iterations. *Journal of Research of the National Bureau of Standards*, 49:33–53, 1952.
- [35] L. Meier III and D. Luenberger. Approximation of linear constant systems. *IEEE Transactions on Automatic Control*, 12(5):585–588, 1967.
- [36] L. A. M. Mello, E. de Sturler, G. H. Paulino, and E. C. N. Silva. Recycling Krylov subspaces for efficient large-scale electrical impedance tomography. *Comput. Methods Appl. Mech. Engrg.*, 199:3101–3110, 2010.
- [37] R. B. Morgan. GMRES with deflated restarting. *SIAM Journal on Scientific Computing*, 24(1):20–37, 2002.
- [38] M. L. Parks, E. de Sturler, G. Mackey, D. D. Johnson, and S. Maiti. Recycling Krylov subspaces for sequences of linear systems. *SIAM Journal on Scientific Computing*, 28(5):1651–1674, 2006.
- [39] T. Penzl. Algorithms for model reduction of large dynamical systems. *Linear Algebra and its Applications*, 415(2–3):322–343, 2006. Special Issue on Order Reduction of Large-Scale Systems.
- [40] A. Ruhe. Rational Krylov algorithms for nonsymmetric eigenvalue problems. II. matrix pairs. *Linear algebra and its Applications*, 197–198:283–295, 1994.
- [41] Y. Saad. *Iterative Methods for Sparse Linear Systems*. 2nd ed., SIAM, Philadelphia, PA, 2003.
- [42] Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869, 1986.
- [43] J. Saak and P. Benner. Efficient numerical solution of the LQR-problem for the heat equation. In *Proceedings in Applied Mathematics and Mechanics*, volume 4, pages 648–649, 2004.
- [44] G. L. G. Sleijpen and H. A. van der Vorst. A Jacobi-Davidson iteration method for linear eigenvalue problems. *SIAM Journal on Matrix Analysis and Applications*, 17(2):401–425, 1996.
- [45] P. Sonneveld. CGS, a fast Lanczos-type solver for nonsymmetric linear systems. *SIAM Journal*

- on Scientific and Statistical Computing*, 10(1):36–52, 1989.
- [46] J. T. Spanos, M. H. Milman, and D. L. Mingori. A new algorithm for L^2 optimal model reduction. *Automatica*, 28(5):897–909, 1992.
 - [47] A. Stathopoulos and K. Orginos. Computing and deflating eigenvalues while solving multiple right-hand side linear systems with an application to quantum chromodynamics. *SIAM Journal on Scientific Computing*, 32(1):439–462, 2010.
 - [48] Z. Strakoš and P. Tichý. On efficient numerical approximation of the bilinear form $c^*A^{-1}b$. *SIAM Journal on Scientific Computing*, 33(2):565–587, 2011.
 - [49] J. van den Eshof and G. L. G. Sleijpen. Accurate conjugate gradient methods for families of shifted systems. *Applied Numerical Mathematics*, 49(1):17–37, 2004.
 - [50] H. A. van der Vorst. Bi-CGSTAB: a fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 13(2):631–644, 1992.
 - [51] H. A. van der Vorst. *Iterative Krylov Methods for Large Linear Systems*. Cambridge University Press, Cambridge, UK, 2003.
 - [52] P. Van Dooren, K. A. Gallivan, and P.-A. Absil. \mathcal{H}_2 -optimal model reduction of MIMO systems. *Applied Mathematics Letters*, 21(12):1267–1273, 2008.
 - [53] S. Wang, E. de Sturler, and G. H. Paulino. Large-scale topology optimization using preconditioned Krylov subspace methods with recycling. *International Journal for Numerical Methods in Engineering*, 69(12):2422–2468, 2006.
 - [54] D. A. Wilson. Optimum solution of model-reduction problem. *Proc. of IEE*, 117(6):1161–1165, 1970.
 - [55] A. Yousuff and R. E. Skelton. Covariance equivalent realizations with application to model reduction of large scale systems. In C. T. Leondes, editor, *Control and Dynamic Systems*, volume 22, pages 273–348. Academic Press, New York, NY, 1985.
 - [56] A. Yousuff, D. A. Wagie, and R. E. Skelton. Linear system approximation via covariance equivalent realizations. *Journal of Mathematical Analysis and Applications*, 106(1):91–115, 1985.